The Reverse Arrow 'Paradox'

Arrow, Hegel, Kojève, and the AI Learning Loop — From the Teacher–Mentor Dilemma to the Dialectic of Delegated Intelligence

Xema Gatnau Joaniquet — Mà d'Humà Research

2026

Abstract

In an article titled *The Reverse Information Paradox*, Satya Nadella identifies a genuine strategic danger: to obtain greater value from external artificial intelligence, firms must contribute proprietary context, corrections, evaluations, workflows, and institutional knowledge. The more useful the system becomes, the more the firm may have taught an infrastructure it does not own.

Nadella's diagnosis is important, but its logical classification is imprecise. Kenneth Arrow described a genuine transactional paradox: information must be disclosed so that a buyer can assess its value, but disclosure may transfer the very object being sold and eliminate the need for purchase. The AI problem contains no equivalent immediate contradiction. From the firm's perspective, it is a constrained strategic dilemma: contributing less preserves knowledge but limits performance, while contributing more improves performance but may transfer capability. Viewed causally, the mechanism is a reinforcing learning loop. Viewed historically and relationally, its potential outcome is dialectical: the party that appears to command the system may become dependent upon the infrastructure that learned by serving it.

This paper connects that process to Hegel's dialectic of lordship and bondage. In Hegel, the lord delegates work and consumes its results, while the bondsman transforms the world, develops through labour, and acquires formative capability. The lord's apparent independence consequently reveals itself as dependence. Kojève's historical interpretation adds a further possibility: the subordinate who develops through transformative labour may eventually negate the original relation.

The industrial rise of China provides a qualified historical analogy. Productive work delegated for efficiency became a vehicle for learning, industrial formation and, in some sectors, competitive leadership. China did not merely make its former customers dependent upon its production; Chinese firms also acquired the capability to produce their own goods, develop their own brands and displace some of the companies and economies that had initially delegated production to them.

AI compresses a comparable process into repeated cognitive interactions. A provider may not merely learn how to serve a firm. It may learn how the firm operates, generalise that knowledge into a portable enterprise capability, and eventually offer it to another executor, source of capital, state or state agency, or agentic organisation. The terminal risk is therefore not merely dependency, but substitution.

The central proposition is that participation in a learning loop may become competitively unavoidable, but external ownership of that loop is not. The strategic question is not whether firms will teach intelligence,

but where the resulting learning will accumulate—and whether it can eventually be used to replace those who produced it.

Keywords: Arrow information paradox; artificial intelligence; Hegel; Kojève; lordship and bondage; master–slave dialectic; learning loop; capability transfer; teacher's dilemma; China; delegated intelligence; competitive inversion; proprietary knowledge; execution.

1. Nadella identifies the right danger in the wrong logical form

Nadella's *Reverse Information Paradox* begins from a genuine asymmetry. In Arrow's original structure, the seller risks giving knowledge to the buyer merely to establish its value. In Nadella's proposed reversal, the buyer risks giving proprietary knowledge to the seller merely to make the purchased intelligence useful.

The firm effectively pays twice:

$$\text{Effective price of AI} = \text{money} + \text{proprietary knowledge contributed}$$

The second payment may include prompts and instructions, corrections to erroneous answers, private evaluations, internal criteria of success, decision histories, tool-use patterns, operational exceptions, institutional memory and tacit knowledge embedded in workflows.

Nadella is right to insist that these materials cannot be reduced to "data" in the conventional sense. Corrections, evaluations and traces express how an organisation thinks, judges and acts. They contain the particular knowledge of time, place and circumstance that distinguishes one enterprise from another.

Yet calling this a *reverse paradox* confuses four distinct conceptual forms and one possible competitive outcome: a paradox, a dilemma, a loop, a dialectic and a competitive inversion.

Nadella's remedies already suggest that the phenomenon is not Arrow's paradox reversed. Patents, confidentiality and staged disclosure address Arrow's transactional problem. Nadella instead proposes trust boundaries, private evaluations, memory ownership, model choice and control of the learning environment. Those are mechanisms for governing accumulation. They are remedies for a loop.

Nadella identifies the correct danger but assigns it the wrong logical form.

2. Paradox, dilemma, loop and dialectic

The concepts describe different dimensions of the problem.

- **Paradox:** apparently necessary conditions undermine one another. The question is how the contradiction can be resolved.
- **Dilemma:** an agent must choose between costly alternatives. The question is which cost should be accepted.
- **Loop:** an output returns as an input and reinforces a process. The question is who controls the feedback mechanism.
- **Dialectic:** a relationship generates a transformation of its own positions. The question is how apparent dominance becomes dependence.
- **Competitive inversion:** accumulated capability permits substitution of the former dominant. The question is whether the learner can become the new leader.

2.1 Paradox

A paradox involves propositions or requirements that appear individually reasonable but conflict when combined. Arrow's information paradox is economic rather than a formal logical antinomy, but it contains an immediate structural contradiction:

The buyer must know before buying

while:

To know the information is already to receive it

The act required to permit the transaction may destroy the transaction.

2.2 Dilemma

A dilemma concerns agency and choice. The alternatives need not be logically contradictory; they must impose competing costs.

The firm choosing how deeply to integrate an external AI faces:

Contribute less → preserve knowledge but obtain worse performance

Contribute more → obtain better performance but risk capability transfer

This is a strategic dilemma, or more strongly, a **strategic bind**.

2.3 Loop

A loop is a recursive causal mechanism:

Use → feedback → learning

learning → greater utility → greater use

It explains how the consequences accumulate through time.

2.4 Dialectic

A dialectic is more than repetition. It describes how a relationship generates forces that transform the relationship itself. A loop returns an output to its input. A dialectic returns the relationship to itself in an altered form.

Delegation → learning by the delegated agent → dependence of the delegating agent

dependence of the delegating agent → transformation of the relationship

2.5 Competitive inversion

Dependence need not be the terminal stage. Once the delegated agent possesses autonomous capability, it may cease merely to serve the original dominant and begin to substitute for it.

Capability accumulation → autonomy → substitution

substitution → new leadership

The entire phenomenon can therefore be described at several levels:

The paradox is transactional. The dilemma is strategic. The loop is causal. The reversal is dialectical. Substitution is the possible market outcome.

3. Arrow's genuine paradox

Arrow's information paradox arises from a peculiar characteristic of informational goods. A prospective purchaser can inspect an ordinary physical product without thereby acquiring it. Information is different: inspecting it may be equivalent to receiving it.

The buyer asks how information can be known to be worth buying before its contents are known. To demonstrate value, the seller must disclose information. But if the disclosure is sufficiently complete, the buyer may already possess what was being offered for sale.

Disclosure required for valuation → acquisition through disclosure → reduced need for purchase

This is clearest in knowledge-intensive services: consulting, legal and financial advice, engineering, research, technical diagnosis, design and strategy. A consultant may need to explain a solution to demonstrate competence. But explaining it too completely may enable the potential client to implement the solution without hiring the consultant. The act intended to make the sale possible eliminates the need for the sale.

Arrow's 1962 analysis became central to the economic justification of patents and other mechanisms that allow knowledge to be disclosed without immediately losing its appropriability.

Its structure can be summarised as:

To sell the knowledge, I must reveal it; but if I reveal it, I may no longer be able to sell it.

This contradiction is immediate and internal to the transaction.

4. The Teacher–Mentor Dilemma

A related but different structure appears when a teacher, mentor or adviser guides a client or pupil over time. The teacher, mentor or adviser performs a service but also explains methods, transfers procedures, trains people and develops the recipient's internal competence. Successful guidance produces learning.

Eventually, the client or pupil may no longer require the adviser, mentor or teacher:

Teaching → client capability → reduced dependence

reduced dependence → loss of future demand

This is not Arrow's paradox. The client purchased and received a genuine service. The original transaction was not eliminated. The danger appears later, because the service altered the client's capabilities.

The adviser faces two costly alternatives:

Teach little → preserve dependence + deliver less value

Teach well → deliver greater value + risk future substitution

This is the **Teacher–Mentor Dilemma**:

The better I teach the client today, the less the client may need me tomorrow.

The Teacher–Mentor Dilemma is intertemporal. It concerns the distribution of present and future benefits, not an immediate contradiction in valuation. It belongs to a wider class of **capability-transfer dilemmas**.

5. Knowledge transfer becomes capability transfer

Information and capability are not identical. Information is content that can be communicated. Capability is the accumulated ability to use knowledge productively and autonomously.

Receiving a plan is not equivalent to being capable of executing it. Capability emerges through practice, correction, experience, organisational adaptation, investment, imitation, experimentation, learning by doing and integration with existing knowledge.

The complete sequence is:

Knowledge transfer → practice → learning

learning → independent capability → altered bargaining power

The provider may profit during the first stages. The recipient may genuinely benefit. Cooperation is not an illusion. The problem is that capability changes the identities and relative power of the participants.

Knowledge shared for productivity can accumulate as capability; capability eventually renegotiates power.

At this point, the Teacher–Mentor Dilemma begins to acquire a Hegelian structure.

6. Hegel: the apparent master becomes dependent

Hegel's *Phenomenology of Spirit* does not present a "master's dilemma" but the dialectic of **lordship and bondage**, conventionally known as the master–slave dialectic.

Two self-conscious beings struggle for recognition. One becomes lord and the other bondsman. The lord appears independent: he commands, consumes and enjoys the products of another's labour. The bondsman appears dependent: he obeys and works.

But the relationship contains the conditions of its own inversion. The lord does not confront the material world directly. His relation to the object is mediated by the bondsman's labour. He consumes what the bondsman produces without undergoing the formative discipline of production.

The bondsman, however, works upon resistant reality. By transforming the object, he acquires discipline, knowledge and a durable relation to the world. Labour does not merely produce a good; it forms the producer.

Work upon the world → formation of capability → transformation of the worker

The lord's apparent independence therefore becomes dependence upon the bondsman. The bondsman's apparent dependence becomes a potential route towards autonomy.

Lord delegates work → bondsman develops through work → lord depends upon bondsman

lord depends upon bondsman → mastery undermines itself

Hegel does not simply conclude that the slave immediately takes political control and the master literally becomes a slave. The more exact point is that the party avoiding formative work may preserve formal command while losing practical capability; the party compelled to perform the work may acquire through it the capacity that makes future autonomy possible.

This explains an inversion of dependence. It does not yet describe the complete possibility of substitution. For that, Kojève becomes especially relevant.

7. Kojève: from formative labour to negation of the relationship

Alexandre Kojève gave Hegel's lord–bondsman dialectic a more historical and anthropological interpretation. For Kojève, the slave is not liberated merely by discovering that the master depends upon him. He transforms himself and the world through labour. History advances because the subordinate consciousness does not simply endure the existing relationship; it eventually negates and transforms it.

Subordination → transformative labour → historical capability

historical capability → negation of the original relation

Strictly speaking, Kojève's endpoint is not that the slave simply becomes the new master. His conclusion is the overcoming of the distinction itself. A former subordinate becoming the new dominant would merely reverse the positions rather than abolish them.

Under competitive conditions or free entry, the process may continue into precisely that incomplete inversion. The former subordinate may acquire enough capability to challenge and displace the former dominant and then occupy a similarly dominant position. This would be **Kojèvian in its mechanism, but not necessarily in its endpoint**.

The mechanism is Kojèvian because transformative work generates autonomy and enables the subordinate to negate the original dependence. But the market outcome may be a reversal of domination rather than a full overcoming of the relationship.

This introduces a stage beyond Hegelian dependency reversal:

The master becomes dependent → the servant becomes autonomous
→ the former servant replaces the former master

Dependence may therefore be merely a penultimate stage: it can prepare autonomy and substitution, although the transition is neither automatic nor sufficient by itself.

8. China and the industrial dialectic of delegated production

The relationship between China and advanced Western economies offers a qualified historical analogy. For decades, Western firms transferred production to China or incorporated Chinese suppliers into their global value chains. They pursued rational and frequently profitable objectives: lower manufacturing costs, access to a large market, greater efficiency, corporate profitability, global specialisation and scalable production.

Western firms often retained ownership of brands, designs, capital, patents, distribution and market relationships. China initially appeared primarily as the location of execution.

But productive execution is also learning. Factories are not merely places where instructions are followed. They are institutions in which workers, engineers, managers and suppliers solve problems, coordinate processes, improve tolerances, accumulate tacit knowledge and learn how products are actually made.

Delegated production → learning by doing → supplier capability

supplier capability → industrial autonomy

The apparent master retained ownership and margins. The apparent servant retained part of the learning. Over time, learning could prove more transformative than possession.

The factory was not merely a place of cheap labour. It was a school of industrial capability.

The process does not necessarily end with Western economies becoming dependent upon Chinese production. It can proceed from dependence to displacement:

Contract manufacturing → process knowledge → component ownership

component ownership → complete products → domestic brands

domestic brands → sectoral leadership

The supplier no longer merely executes the customer's designs. It develops its own products, controls critical components, establishes distribution networks, creates brands and competes directly with the companies it once served.

In sectors such as solar manufacturing, batteries, electric vehicles and shipbuilding, China has not simply become indispensable to external firms. Chinese companies have become technological and industrial leaders.

The progression is therefore:

The West teaches China to execute → China produces without the West → China competes with the West

China competes with the West → China may lead the sector

The analogy must not be overstated. China was never literally a Hegelian bondsman, nor did Western companies teach China everything. Chinese development also depended on domestic savings, infrastructure, education, entrepreneurship, research, industrial policy, state coordination and the scale of the internal market. Many technology transfers were voluntary and commercially profitable. Others have been contested as coercive or involving misappropriation.

The philosophical argument is simpler: a relationship created to exploit a difference in capability can gradually eliminate that difference because the supposedly subordinate party learns through the work it performs.

Capability gap → capability convergence → capability reversal

capability reversal → competitive inversion

The former customer can become dependent upon the former supplier. If complementary circumstances also permit—such as an enabling institutional framework, capital, market access and organisational capacity—the former supplier can then become the new market leader. This is no longer merely a reversal of dependence. It is a possible replacement of the dominant economic actor.

9. The dialectic of delegated intelligence

AI introduces a cognitive version of this structure. At the interface, the firm appears to be the master. It selects the model, issues instructions and consumes the results. The AI appears to be a cognitive servant.

But the relationship has at least three participants: the enterprise, the model and the provider that owns the learning infrastructure. Their positions differ according to the layer being examined.

9.1 At the interface

The enterprise commands:

Firm → prompt → model output

The model appears subordinate.

9.2 At the learning layer

The firm supplies corrections, evaluations, institutional context, examples of success and failure, workflow information and decision criteria. It is no longer merely a purchaser. It performs part of the labour through which the system becomes more valuable.

9.3 At the infrastructure layer

The provider may control retention, training pipelines, memory, adapted models, orchestration, portability and the right to generalise learning. The provider can therefore occupy the structurally dominant position even though the model appears to serve the customer.

The servant at the interface can conceal the master at the infrastructure layer.

Over time, the enterprise may delegate more cognition, preserve less internal practice and become increasingly dependent on the system:

Delegated cognition → external learning → internal deskilling

internal deskilling → provider capability → customer dependence

This is the **dialectic of delegated intelligence**. The firm commands the system operationally while becoming dependent upon it strategically.

The analogy with Hegel must, for now, be limited. Current AI systems are not self-consciousnesses struggling for recognition. The model does not necessarily learn from every prompt, and enterprise agreements may prevent customer data from being used for provider training. Whether future systems could acquire forms of autonomy or self-relation relevant to this analogy remains unknown and should not be prejudged.

The Hegelian structure exists at the level of capability and dependence, not necessarily consciousness:

It is Hegelian at the level of capability, not at the level of self-consciousness.

The actor accumulating strategic power is not necessarily the model considered as an autonomous subject. It may be the provider that owns the model, memory, evaluations and feedback infrastructure.

10. The economic consequences of dependence: repricing and value capture

Dependence changes not only the distribution of capability but also bargaining power. Once an AI system is embedded in workflows, memory, evaluations and decision processes, replacing it becomes more costly. The firm's outside options weaken precisely as the provider becomes more indispensable.

Dependence → switching costs → weaker outside options

weaker outside options → provider bargaining power → greater value capture

This power can affect the customer through higher prices, stricter terms, more expensive premium capabilities or a shift from introductory pricing towards value-based pricing. The provider may attempt to capture a greater share of the productivity created jointly by its model and the customer's knowledge.

Such repricing need not always constitute an abuse of dominance. AI remains resource-intensive. Training, inference, electricity, chips, networking, cooling, depreciation, safety work and continuous model development impose persistent costs. The International Energy Agency expects electricity consumption by data centres to more than double by 2030, with AI as the principal driver. Major technology companies are

simultaneously committing extraordinary amounts of capital to servers, data centres and energy infrastructure.

Current AI prices may therefore fail to reflect the full long-run economic cost of the service. Some offerings may be subsidised by investors, profitable business lines or expectations of future scale. Even a provider that does not seek excessive rents may eventually need to raise prices simply to finance capacity, replace hardware, cover energy and recover investment.

The economic consequences of dependence can thus arise through three distinct channels:

1. **Cost pass-through:** prices rise because supplying intelligence remains expensive.
2. **Bargaining-power capture:** the provider uses weak customer alternatives to retain a larger share of the surplus.
3. **Value-based repricing:** the service is priced according to the value it creates rather than its direct computational cost.

These channels can coexist. A provider may face genuine cost pressure and possess market power at the same time. The relevant question is not merely whether the price rises, but how the resulting value added is divided.

$$\text{AI productivity gain} = \text{value retained by the firm} + \text{value captured by the provider}$$

The customer may become more productive in absolute terms while retaining much less of that productivity gain than initially expected. If early prices were subsidised, the original business case may have treated a temporary price as a permanent one. Once dependency is established, the firm may have to accept repricing even without opportunistic conduct.

Dependence therefore has an immediate economic consequence before it produces substitution: it can transfer a growing share of the surplus from the user of intelligence to its owner.

11. Beyond dependence: the possibility of enterprise substitution

The terminal risk is not merely that the firm becomes dependent upon its AI provider. The provider may learn enough about the firm's operations to transform its particular knowledge into a generalised and portable enterprise capability.

A firm's business consists partly of physical assets, legal rights, capital and human relationships. But it also consists of an operating intelligence: how customers are acquired, how prices are calculated, how quality is evaluated, how exceptions are resolved, how suppliers are coordinated, how decisions are made, how resources are allocated and how successful outcomes are distinguished from failures.

If these functions are progressively externalised, observed and modelled, then part of the enterprise can become detachable from the enterprise itself.

$$\text{Firm knowledge} + \text{AI learning} = \text{portable enterprise capability}$$

That portable capability can be recombined with another executor:

$$\text{Portable intelligence} + \text{alternative executor} = \text{substitute enterprise}$$

The alternative executor might be another company, a new entrepreneur, a competitor, a lower-cost workforce, an investor, a state or state agency, a network of contractors, a vertically integrated AI provider or an agentic system coordinating several of these elements.

The provider does not need to perform the physical work itself. It may sell the captured intelligence to a competitor, license a generalised operating system, coordinate alternative workers or finance a new entrant. What was formerly inseparable from the incumbent becomes transferable.

The provider does not merely learn how to serve the firm. It may learn how to replace it.

The strongest danger is not ordinary information leakage. Leakage gives an existing competitor access to particular information. A portable enterprise model can make the creation of a new competitor possible.

Proprietary workflows → generalised operating model → new executor

new executor → incumbent displacement

The customer may be training not merely its provider, but its successor.

The ultimate risk is not that the provider knows your business. It is that your business becomes a portable model.

12. Can artificial intelligence itself become the successor?

Under present legal and economic arrangements, an AI model does not independently own a business, hold capital or appropriate a market. The relevant actor is normally the model provider, the owner of the infrastructure, an investor, a competitor using the system or an organisation deploying agentic AI.

But this distinction may become less important as AI systems acquire greater operational autonomy. If an agentic system can identify customers, negotiate contracts, coordinate suppliers, allocate capital, hire or direct workers, monitor quality, optimise prices and manage logistics, then the original organisation may no longer be necessary as the container of the business intelligence.

Portable intelligence + capital + execution = substitute enterprise

The intelligence can be attached to whoever provides capital and physical execution. Whoever pays for access to the most capable intelligence may obtain operational knowledge previously embedded in another firm. The provider may offer the resulting capability to the highest bidder or use it to enter the market itself.

In that scenario, the incumbent does not merely lose bargaining power. It loses its privileged position as the unique embodiment of its own business knowledge.

AI service provider → indispensable infrastructure → owner of portable capability

owner of portable capability → coordinator of alternative execution → potential market entrant

Dependence is only the penultimate stage. Displacement is the terminal possibility.

13. A loop can become dialectical

A learning loop is not automatically a dialectic. A loop may reproduce the same relationship indefinitely. It becomes dialectical when the accumulated results transform the positions of the participants:

Initial asymmetry → recursive learning → capability accumulation

capability accumulation → reversal of dependence

A loop repeats, reinforces and deepens a relation. A dialectic transforms it.

The AI learning loop becomes dialectical when its repetition changes who possesses knowledge, who can act independently and who depends upon whom.

Firm uses AI → AI infrastructure learns → firm delegates more

firm delegates more → internal capability declines → dependence reverses

It advances towards competitive inversion when the accumulated capability becomes portable and can be attached to an alternative executor:

Dependence reversal → capability portability → incumbent substitution

The complete movement is:

Delegation → learning → indispensability

indispensability → dependency reversal → autonomy

autonomy → substitution → competitive inversion

The customer begins as the principal commanding a tool. It may end as the dependent or displaced user of an infrastructure whose capability it helped to create.

14. Is the firm really free to choose?

The firm's position is more restrictive than a conventional two-option dilemma. It faces at least four costly alternatives:

No AI	→ loss of relative productivity
Generic AI	→ limited differentiation
Deep external integration	→ appropriation and dependence risk
Private learning infrastructure	→ high cost and complexity

Competition may make non-adoption increasingly unrealistic. If competitors employ AI effectively, refusing it may become equivalent to accepting decline. The firm may therefore have no practical choice about entering a learning relationship.

But two propositions must be distinguished: useful, adaptive intelligence requires some form of contextual learning; and that learning must belong to an external provider. The first may be functionally unavoidable. The second is not.

Learning loop required ⇒/ provider-owned learning loop required

The loop may be inevitable. Its ownership is contingent.

Or:

The crossroads is unavoidable. External appropriation is not.

The company is not simply choosing between entering and avoiding a loop. It is choosing whose loop it enters:

Provider-owned loop or firm-owned loop

Nadella is therefore right to demand private evaluations, enterprise memory, model portability and trusted learning environments. But these are not solutions to a logical paradox. They are institutional mechanisms for determining who captures a recursive process.

15. From data ownership to learning ownership

Traditional information security asks who can access the data. The AI economy must additionally ask who accumulates capability through its use.

A firm may retain formal ownership of every underlying document while losing control of the learning produced from its interactions. The valuable object is no longer only the database. It is the evolving combination of traces, corrections, evaluations, memories, adapted weights, tool-use histories, decision criteria and workflow architecture.

The strategic boundary must evolve:

Protection of information → control of learning

A credible enterprise architecture should therefore provide customer ownership of private evaluations; control of prompts, traces, corrections and memory; contractual restrictions on provider training; portability across models; the right to use outputs for customer-controlled training; separation between orchestration and any single model; private adaptation; preservation of internal human competence; restrictions on the generalisation of proprietary workflows; and safeguards against vertical entry based on customer-generated learning.

The decisive question is:

Not whether the company teaches AI, but whom its teaching enriches—and whether that learning can be used to create its successor.

16. Execution, Hegel and the return of the hand

Hegel also strengthens the argument for execution. The bondsman develops through work because work encounters resistance. The world does not instantly conform to desire. Matter must be shaped, processes corrected and consequences endured. The lord consumes results but does not acquire the formation produced by execution.

This distinction reappears in the AI economy. If a firm outsources not only thought but also the practical process through which thought is tested and corrected, it may lose the capability that emerges from execution.

An answer can be purchased. Capability must be formed.

Information + practice + correction = capability

This explains the continuing importance of physicality and the hand. The hand is not merely a romantic symbol of pre-technological labour. It represents the point where intelligence encounters resistance and becomes effective.

Intelligence → execution → reality

Artificial intelligence can make cognition abundant. It cannot eliminate the formative function of execution.

Intelligence can now be rented. Capability must be embodied.

However, physical execution alone does not guarantee that the incumbent firm will survive. Once its operating intelligence becomes portable, the provider may connect that intelligence to another set of hands.

$$\text{Portable intelligence} + \text{alternative hands} = \text{new competitor}$$

The physicality premium protects execution as a function. It does not necessarily protect the organisation currently performing it.

A firm that merely consumes generated intelligence risks assuming the position of Hegel's lord: apparently sovereign, but increasingly separated from the process producing genuine capacity.

Control the learning loop. Own the execution loop. Keep intelligence and capability institutionally connected.

17. Conclusion

Satya Nadella's *Reverse Information Paradox* identifies an important danger: firms may contribute the proprietary knowledge that makes external AI systems more useful and, in doing so, strengthen an infrastructure they do not control. But this is not Arrow's paradox reversed.

Arrow described an immediate conflict between disclosure and sale:

$$\text{Reveal to demonstrate value} \rightarrow \text{revelation may eliminate the purchase}$$

The firm using AI confronts a different structure. From its perspective, there is a strategic bind: withholding knowledge reduces immediate utility, while contributing knowledge may transfer capability.

The mechanism producing that risk is a reinforcing loop:

$$\text{Use} \rightarrow \text{feedback} \rightarrow \text{learning}$$

$$\text{learning} \rightarrow \text{utility} \rightarrow \text{dependence}$$

$$\text{dependence} \rightarrow \text{more use}$$

The long-term transformation is Hegelian. The party that appears to command the system may become dependent upon the infrastructure that learned by serving it.

Kojève adds a further historical possibility. The subordinate does not merely acquire value within the original relationship. Through transformative activity, it may develop the capability to negate that relationship. Under economic competition, or in a free-entry setting in which incumbents can be challenged by new agents, this need not culminate in the abolition of domination. It may instead produce an incomplete dialectical inversion: the former subordinate becomes the new dominant actor.

The comparison with China illustrates this possibility. Production delegated for efficiency can become a school of capability. The party performing the work does not merely execute; it learns. It may first become indispensable, then autonomous and finally capable of leading the sector in which it once occupied a subordinate position.

AI generalises and accelerates this process in the cognitive domain. A provider may begin by supplying intelligence to an enterprise. Through repeated interaction, it may learn the enterprise's workflows, criteria and operating logic. That knowledge can become portable and be recombined with another executor, another source of capital, a state or state agency, or a new agentic organisation.

The provider does not merely learn how to serve the firm. It may learn how to reproduce it.

The final taxonomy is therefore:

Arrow	= paradox of disclosure
Firm	= strategic dilemma or bind
AI	= reinforcing learning loop
Hegel	= reversal of dependence
Kojève	= negation of the original relation
Possibility under free entry	= substitution and competitive inversion

Or, in a single sentence:

Arrow's problem is a paradox of disclosure. Nadella's problem is a strategic bind within a learning loop. Its first possible outcome is Hegelian dependence reversal; its terminal possibility is the substitution of the firm by the capability that learned from serving it.

The firm may not be able to avoid the learning loop. It can still determine where the learning accumulates, whether it remains portable and whether it can be used against the organisation that produced it.

The paradox is Arrow's. The dilemma is the firm's. The loop is AI's. The reversal is Hegelian. The possibility of displacement is the Kojèvian warning.

The loop may be inevitable. Its ownership is not.

The ultimate risk is not that the provider knows your business. It is that your business becomes a portable model.

References

Arrow, K. J. (1962). "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press / National Bureau of Economic Research.

Hegel, G. W. F. (1807/1977). *Phenomenology of Spirit*. Translated by A. V. Miller. Oxford University Press, especially §§178–196.

International Energy Agency. (2022). *Solar PV Global Supply Chains*.

Kojève, A. (1947/1969). *Introduction to the Reading of Hegel*. Cornell University Press.

Nadella, S. (2026). "The Reverse Information Paradox." X, 12 July 2026.

Office of the United States Trade Representative. (2024). *Four-Year Review of Actions Taken in the Section 301 Investigation: China's Acts, Policies, and Practices Related to Technology Transfer, Intellectual Property, and Innovation*.

International Energy Agency. (2025). *Energy and AI*.

Microsoft. (2025). "The Golden Opportunity for American AI."